REPORT RESUMES

ED 011 047 56

AN EXAMINATION OF HUMAN STRATEGIES FOR ACQUIRING INFORMATION. PSYCHOLOGICAL AND EDUCATIONAL FACTORS IN TRANSFER OF TRAINING, TECHNICAL REPORT NO. 8.

BY- DAVIS, DANIEL J. STOLUROW, LAWRENCE M.
ILLINOIS UNIV., URBANA, BUR. OF EDUC. RESEARCH

REPORT NUMBER BR-5-0230-8 PUB DATE OCT 65
REPORT NUMBER NDEA-VIIA-C-1288-8
CONTRACT OEC-4-20-002
EDRS PRICE MF-$0.09 HC-$1.96 49P.

DESCRIPTORS- *RESEARCH METHODOLOGY, *STIMULUS BEHAVIOR, *RESPONSE MODE, *QUESTIONING TECHNIQUES, EXPERIMENTS, *INFORMATION PROCESSING, LEARNING PROCESSES, URBANA

AN EXPERIMENTAL SITUATION WAS DEVELOPED IN WHICH THE SUBJECT HAD THE TASK OF DETERMINING WHICH OF A SET OF POSSIBLE EVENTS HAD OCCURRED ON EACH TRIAL. THIS WAS DONE BY ASKING THE EXPERIMENTER QUESTIONS WHICH COULD BE ANSWERED "YES" OR "NO" UNTIL THE EVENT WAS NAMED. FOUR EXPERIMENTS WERE CONDUCTED IN WHICH STRATEGIES WERE EXAMINED AS A FUNCTION OF VARIOUS ENVIRONMENTAL FACTORS. THE OVERALL RESULTS OF THE EXPERIMENTS INDICATED THE THREE MAIN FACTORS THAT INFLUENCED THE USE OF STRATEGIES WERE (1) THE EXTENT TO WHICH THE QUESTIONS REFLECTED THE DOMINANT CHARACTERISTICS OF THE STIMULI, (2) THE AVERAGE AMOUNT OF INFORMATION WHICH WAS OBTAINED WITH QUESTIONS, AND (3) THE RISK OF HAVING TO USE A LARGE NUMBER OF QUESTIONS. (TC)

C-1288
5-0230

TRL

TRAINING RESEARCH LABORATORY

Department of Psychology | Bureau of Educational Research

University of Illinois | 8 Lincoln Hall | Urbana, Illinois

AN EXAMINATION OF HUMAN STRATEGIES FOR ACQUIRING INFORMATION

Daniel Jay Davis

Technical Report No. 8

October 1965

Psychological and Educational Factors in Transfer of Training

U. S. Office of Education
Contract OE 4-20-002

Lawrence M. Stolurow
Principal Investigator

TRAINING RESEARCH LABORATORY
University of Illinois
Urbana, Illinois

AN EXAMINATION OF HUMAN STRATEGIES FOR ACQUIRING INFORMATION[1,2]

Daniel Jay Davis

Technical Report No. 8

October 1965

Psychological and Educational Factors in Transfer of Training
U. S. Department of Health, Education and Welfare
Office of Education
Title VII Contract OE 4-20-002

Lawrence M. Stolurow
Principal Investigator

[1]Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Psychology in the Graduate College of the University of Illinois, 1965.

[2]Research was also supported in part by the Office of Naval Research under contract US NONR 3985(04), Lawrence M. Stolurow, Principal Investigator.

AN EXAMINATION OF HUMAN STRATEGIES FOR ACQUIRING INFORMATION

Daniel Jay Davis, Ph. D.

Department of Psychology

University of Illinois, 1965

An experimental situation was developed in which the subject (S) had the task of determining which of a set of possible events had occurred on each trial. This was done by asking the experimenter (E) questions which could be answered "yes" or "no" until the event was named. An event consisted of one of a set of cards being "correct." Each card was labeled in some way (e.g., A, B, ---) and had a fraction written on it indicating its probability of being correct. The correct card was determined on each trial according to a random process in which outcomes had the corresponding probabilities of the cards.

After a fixed number of trials, S was asked to describe what he considered to be the best single way of asking his questions. This strategy was the primary datum of the experiment. A measure of "goodness" was determined which was the average number of questions which would be required per trial if the strategy were used over a long series of trials.

The experimental situation was viewed as a communication system in which the set of cards and the random process which generated outcomes (correct cards) determined an information source. The experimenter behaved as a channel which had the capacity of one bit per question and S's strategy encoded each outcome into a sequence of "yes" and "no" answers. As a result, it was possible to determine the lower bound on the average number of questions required for a given source and the efficiencies of strategies could be compared.

Four experiments were conducted in which strategies were examined as a function of various environmental factors.

In Experiment I it was found that a question based on the dominant characteristics of the stimuli was preferred as long as it was not highly inefficient. When the question became highly inefficient, it was dropped in favor of less dominant but more efficient questions.

In Experiment II it was found that efficiency generally decreased as the skewness of the source distribution increased. This was a result of the tendency of Ss to ask questions which halved the cards in terms of number rather than probability. It was inferred from comments by Ss that this tendency resulted from a concern with the risk of asking a large number of questions.

In Experiment III, it was observed that experience generally led to the development of more efficient strategies. This indicates that the Ss were adaptive in the sense that they learned to pay less for each bit of information obtained from the environment. However, in some cases, efficiency was sacrificed to keep the risk of using a large number of questions low. In Experiment IV, a change in approach to one situation was reflected in a change to similar situations. This provides evidence that strategies involved the use of rules for applying the three factors which generalized across situations.

On the basis of the results of the experiments, it was concluded that three main factors influenced the use of strategies. They were: (1) the extent to which the questions reflected the dominant characteristics of the stimuli; (2) the average amount of information which was obtained with questions; (3) the risk of having to use a large number of questions.

PRECEDING PAGE BLANK- NOT FILMED

TABLE OF CONTENTS

INTRODUCTION

A common task faced by humans is that of transforming an uncertain situation into one which is less uncertain. Generally this is accomplished by probing or asking a question; observing a result or reply; and continuing the process until the desired goal is achieved.

An example of this task is electronic "troubleshooting" in which a technician must determine the source of a malfunction for a piece of equipment. His procedure is to apply successive tests until all possibilities but one are eliminated.

The purpose of this research program is to examine the behavior of adult, normal, human subjects in the above type of situation. The examination focuses on the particular sequence of questions or probes chosen by a subject and on the factors which influence the choice. In the remainder of this paper the word "strategy" is used in place of "sequence of questions."

The program consists of two main parts:

1. The development of a general experimental situation in which:
   a. Strategies can be observed as a function of independent experimental variables.
   b. Strategies can be described and compared in terms of a quantitative measure of "goodness"(to be defined later).
2. The design and execution of experiments which illustrate the influence of various environmental factors on the use of strategies.

The Experimental Situation

In the experimental situation the subject (S) has the task of determining which one of a set of possible events has occurred. He does this by asking the experimenter (E) questions of a prescribed type (e.g., binary) and by using the answers to reduce the set of possibilities until the event which occurred is named. The ordered set of questions asked by S defines his strategy.

In order to derive a measure of "goodness" for strategies, two restrictions must be placed on the experimental situation. First, the set of possible events must be clearly defined. If S's set of possible events is not identical to E's set, then S could employ a strategy which is "good" in terms of his set but not in terms of E's set. Second, each event must have a well-defined probability of occurrence. If S considers certain events as likely to occur which E has made unlikely to occur, then S's strategy could be "good" under his assumptions but not under E's.

With these restrictions in mind, the following situation was devised. Both E and S are seated at a table, and a set of cards is placed before S. Each card is labeled in some way (e.g., A, B---) and has a fraction written on it. The E explains that they are going to go through a series of trials, and on each trial one of the cards will be considered correct. The fraction on a card indicates its probability of being correct on each trial. The meaning of the probability is explained to S: "A probability of 1/4 means that the card will be correct about once out of every four trials." In addition, S is told that the outcome on a given trial does not affect outcome probabilities on following trials and that there are no cycles on the outcomes.

The S is shown the device which determines the correct card -- a circular dial with a rotatable pointer in the center. The circumference is divided into regions, each of which corresponds to one of the cards. The arc-length of a region is proportional to the probability of occurrence of the corresponding card.

On each trial, E spins the pointer which is now out of S's view, and uses the outcome to determine the correct card. The S must determine the correct card by asking questions which can be answered "yes" or "no." For example, a typical question would be: "Is the correct card A, B, or C?" E would answer "yes" if the correct card were A, B, or C, and "no" otherwise. At the end of a fixed number of trials, S is asked to describe what he considers to be the best single way of asking his questions, i. e., the best first question; the best second question, depending on how the first question was answered, etc. This final strategy is the primary datum of the experiment.

A Proposed Measure of "Goodness"

For a given set of cards and associated probabilities, it is possible to calculate the average number of binary questions needed to determine the correct card when a fixed strategy is used over a long series of trials. Consider the case with n cards where the probability that the ith card is correct is given by p(i), i=1,---n. Let $m(\sigma,i)$ be the number of questions required when card i is correct and strategy σ is used. It follows that the average number of questions required with strategy σ is given by equation 1.

$$(1) \qquad M(\sigma) = \sum_{i=1}^{n} p(i)\, m(\sigma,i)$$

As an example, Table 1 illustrates a case with six cards, and Figure 1 shows the flow diagrams which result when two strategies, σ_1 and σ_2, are used with this case. Using equation (1), $M(\sigma_1)$ and $M(\sigma_2)$ are calculated as follows:

$$M(\sigma_1) = 1/2\ (2) + 1/4\ (2) + 1/4\ (3) = 2.25$$

$$M(\sigma_2) = 1/2\ (1) + 1/4\ (2) + 1/4\ (4) = 2$$

Thus, σ_1 requires 2.25 questions on the average to determine the correct category while σ_2 requires 2 questions.

It is proposed to use $M(\sigma)$ as an index of the "goodness" of strategies. A low value of $M(\sigma)$ indicates that fewer questions are needed on the average to determine the correct card. In this sense, σ_2 is better than σ_1 since it would be cheaper to use if each question cost money.

In some cases it is desirable to have a measure of efficiency for strategies. A useful quantity for this purpose is the lower bound on $M(\sigma)$. The lower bound (L) can be determined by using Shannon's 9th theorem which is the fundamental theorem for a noiseless channel (Shannon and Weaver, 1949). It states that if an information source has an entropy

Table 1

A Situation With Six Cards

Card	1	2	3	4	5	6
Probability	1/2	1/4	1/16	1/16	1/16	1/16

or uncertainty U (bits per symbol) and a channel has a capacity C (bits per second), then it is not possible to transmit at an average rate greater than C/U symbols per second.

In terms of the experimental situation described here, the set of cards and the random process which generates outcomes (correct cards) determine an information source. If the outcome probabilities on any trial are independent of outcomes on previous trials, the uncertainty of the source is given by equation 2.

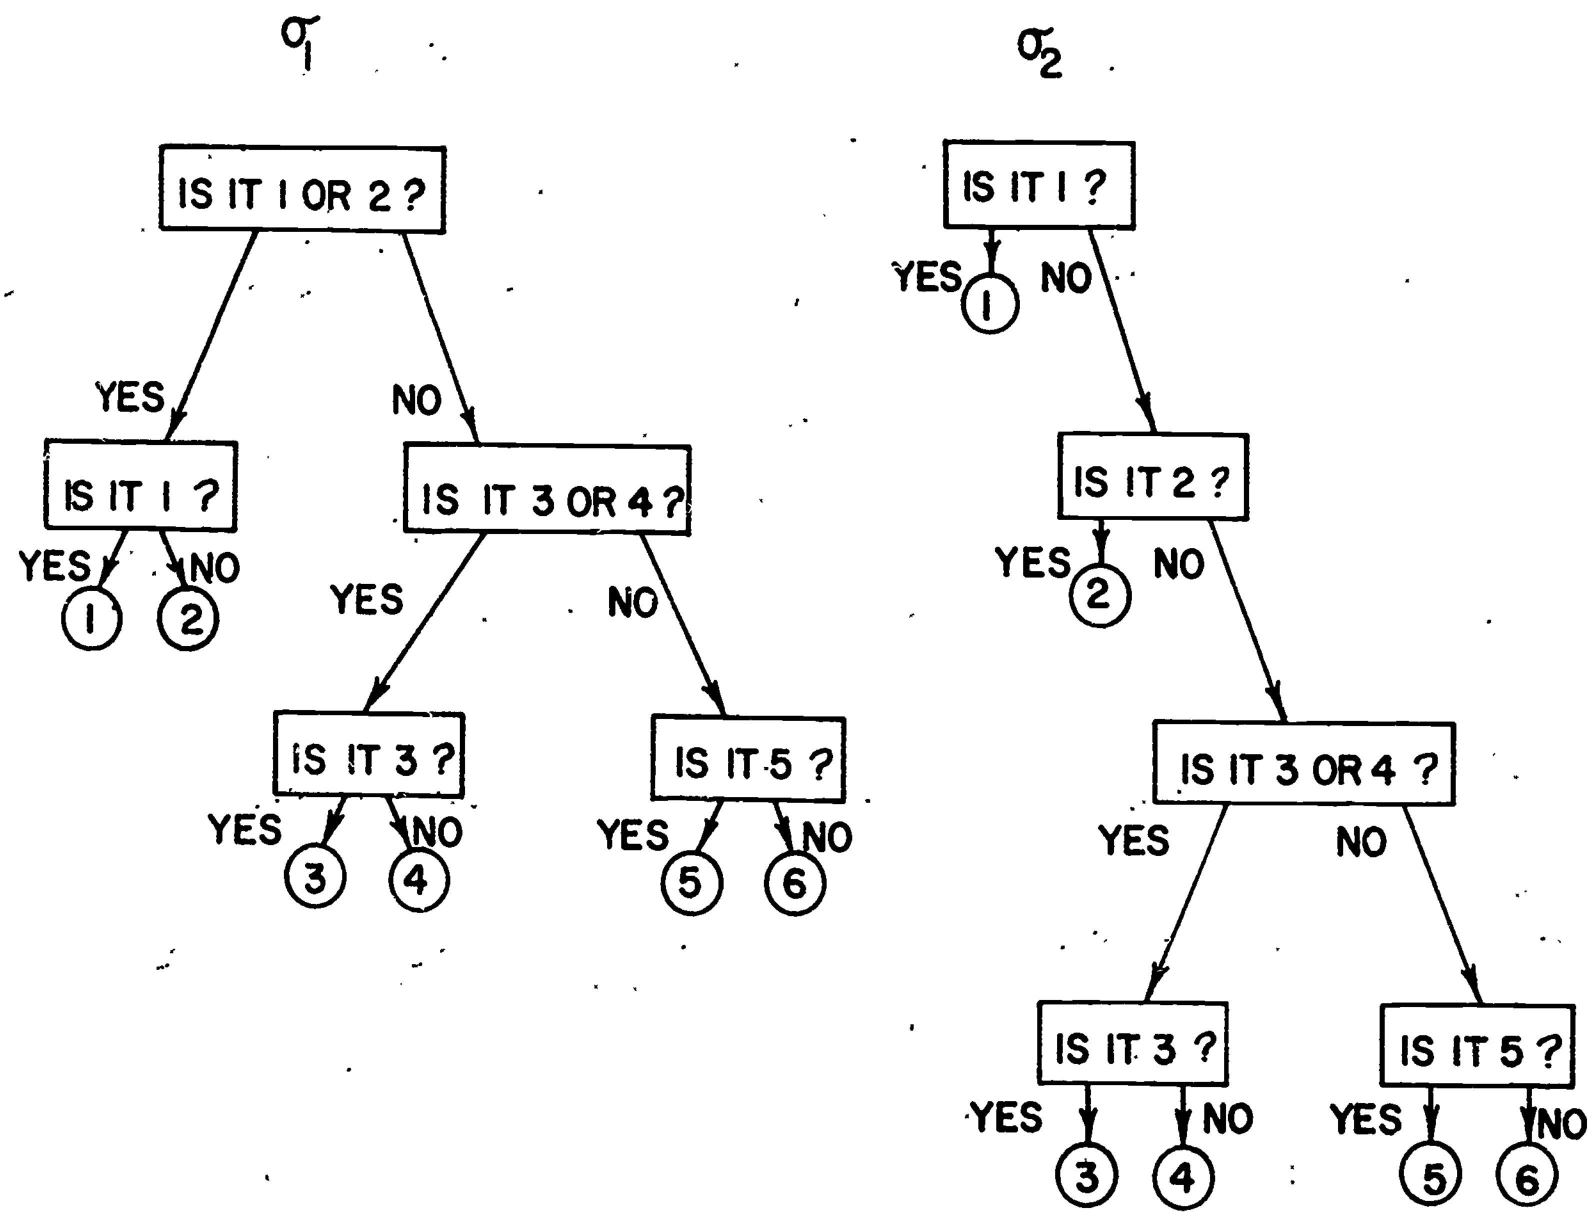

Figure 1. Flow diagrams of strategies σ_1 and σ_2 applied to the case of Table 1.

(2) $$U = \sum_{i=1}^{n} p(i) \log_2 p(i) \text{ bits per card}$$

The experimenter behaves as a channel which transmits one of two symbols ("yes" and "no") after each question. The capacity of the channel is then 1 bit per question.

The theorem applied to this situation states that 1/U cards per question is the maximum average rate of transmitting information concerning the set of cards. An equivalent statement is that U questions per card is the minimum average number of questions needed to determine the correct card. This is L - the lower bound on $M(\sigma)$.

The uncertainty of the source described in Table 1 is 2 bits per card as calculated with equation 2. This means that no strategy can be found which uses less than 2 questions on the average in determining the correct card. This lower bound is achieved by σ_2 and therefore it is a best strategy in terms of $M(\sigma)$.

The quantity $[M(\sigma) - L]$ is a measure of inefficiency and is referred to as the "surplus cost" of a strategy. It is the average number of questions required on each trial above what is absolutely necessary. For example, in the case of Table 1, the surplus cost of σ_1 is .25 questions per trial, while the surplus cost of σ_2 is zero questions per trial. If a dollar is required for each question and S is given credit for two dollars at the beginning of each trial, the use of σ_1 results in an average loss of $.25 per trial, while the use of σ_2 results in zero average loss. Thus, the surplus cost can be translated into a measure of extra expense in terms of money, energy, time, or whatever S must pay for each question.

In the following experiments, the source probabilities have been chosen so that at least one strategy exists which achieves the lower bound of $M(\sigma)$. Such a strategy is optimal and uses questions which divide the set of possible cards into equiprobable groups.

Review of the Literature

There have been relatively few reports in the psychological literature of experiments dealing with strategies for acquiring information. One reason for this is that a strategy is a complex unit of behavior which is difficult to describe quantitatively. Also, in even the simplest of situations, the number of possible strategies is extremely high.

In view of these difficulties most applications of information theory in psychology have been in experiments where *S* is a passive receiver of information (Garner, 1962; Posner, 1964). For example, in his information analysis of the game "Twenty Questions" Bendig (1953) did not allow his *Ss* to ask questions freely. Rather, they heard questions asked and answered and then tried to guess the correct topic. After each question the uncertainty of the topics was estimated using the proportion of *Ss* selecting each of the possible topics. Although this experiment illustrated how effective *Ss* were in using information provided by specific questions, it did not touch on the ability of *Ss* to ask questions to obtain information.

Rimoldi (1963) attempted to overcome the difficulties of dealing with strategies by giving *Ss* a list of questions to use. In order to solve the task, *Ss* chose questions which were then answered by *E*, and the sequence of questions chosen defined *S*'s strategy. For each question in a given situation a "utility index" was empirically determined as the proportion of *Ss* who used the question. It was assumed that the more a question was used by *Ss*, the greater was its usefulness in solving the task. A "good" strategy was defined as one which used questions with a high utility index first.

There are two difficulties with Rimoldi's approach. In the first place, the use of a list of questions greatly restricts and influences S's choice of strategy. In the second place, the use of a question by Ss is not necessarily a reflection of its usefulness in obtaining information. There is evidence that questions are also chosen on the basis of their ease of use in the task situation. This point is discussed in Experiment I.

Glaser, et. al. (1954) and Glaser and Schwarz (1954) described an approach which is similar to Rimoldi's. The task was to determine a fault in a system by observing the results of various diagnostic tests. A list of tests was provided and S chose to see the results of some of these. When S felt that he had obtained enough information, he chose the most likely fault from a list of possible faults. It was assumed that S considered the faults to be equiprobable before seeing the results of the diagnostic tests. Therefore, the initial uncertainty of the set of possible faults was $\log_2 n$ bits when there were n possible faults. If S chose the correct answer on the first try, he was assumed to have obtained $\log_2 n$ bits of information from the diagnostic tests. If he chose the answer on a later try, the amount of information obtained was lower. The efficiency of a strategy was the amount of information obtained per diagnostic test. Unfortunately, no empirical findings were reported for this task.

Detambel (1956), Detambel and Stolurow (1957), Goldbeck, et. al. (1957), and Dale (1959) reported experiments in which Ss had to detect a fault in a system. In order to simplify the situation, Detambel and Stolurow used a task in which the number of possible strategies was small. Dale and Goldbeck, et. al. restricted their examinations to cases in which the possible outcomes were equiprobable. These last authors were interested in S's use of the "half-split" technique where each probe or question divides the set of remaining possibilities into two equal groups. The half-split technique is the optimal approach when all possible outcomes have equal probability.

Several reports dealt with procedures for determining optimal search strategies in complex systems (Gluss, 1959; Stolurow, et. al., 1955). However, these were not directly concerned with human performance in the situation. The criterion of optimization in the two cited reports was the average work required to detect a fault. Gluss was concerned with the problem in the general case, while Stolurow, et. al., used maintenance data to determine optimal trouble shooting procedures for aircraft.

Pertinent results of the above studies are mentioned in the following chapters. It is felt that the approach of this report goes beyond the approaches cited above in that a meaningful quantitative index of "goodness" can be determined for strategies and compared for different situations. Also, Ss are not greatly restricted in their use of questions and the large number of possible strategies does not constitute a barrier to experimental investigation.

EXPERIMENT I: THE INFLUENCE OF LABEL CHARACTERISTICS ON THE USE OF STRATEGIES

Several experimenters have observed that Ss generally do not use highly efficient strategies for acquiring information (Detambel and Stolurow, 1957; Dale, 1959). This suggests that factors other than efficiency enter into determining Ss choice of a question, and one such factor seems to be the way in which it reflects the dominant or readily apparent characteristics of the task stimuli (Bruner, et. al., 1957).

Questions which partition the task stimuli along dominant dimensions are preferred to questions which do not. For example, if the stimuli are easily distinguished in terms of one dimension (say color), but not in terms of another (say size), then questions which are stated in terms of color might be preferred even if they are less efficient than questions stated in terms of size. In this way, efficiency is sacrificed in order to use questions which are based on the dominant characteristics of the stimuli.

The purpose of this experiment was to examine the relative influence on strategies of question efficiency and question dominance in terms of the stimulus structure. In the experimental situation of this report it was observed that Ss tended to phrase their questions in terms of dominant characteristics of the card labels. If the cards were labeled with numbers, the most common questions were, "Is the correct card even?" and "Is the correct card odd?" If half the cards had a letter on them and half had a number, the question, "Is the correct card a letter?" was almost always the first question.

An interesting situation arises when a dominant label distinction leads to the use of an inefficient question. For 16 equally likely cards, this would be the case if there were nine cards with a letter and seven with a number. The question, "Is the correct card a letter?" is still the dominant question, but is now slightly inefficient. If there are 13 cards with a letter and three cards with a number, this question is highly inefficient. Thus, the efficiency of a question can be varied depending upon how it partitions the stimuli.

It was expected that Ss would persist in using a question based on a dominant label distinction as long as it was not highly inefficient. When such a question was highly inefficient it would be dropped in favor of another which was less dominant but more efficient.

Method

Experimental design. Eight treatment groups were used. Each group consisted of nine males and nine females. In all treatments a source of 16 equally likely cards was used, and each card was labeled with either a letter or a number. The number of cards labeled with a letter was different for each treatment and ranged from 8 to 15. The set of labels are shown in Table 2.

In each case the dominant label distinction was the letter-number dichotomy. However, the amount of information obtained by using a first question based on this distinction varied from one bit at eight letters to about .34 bits at 15 letters. In the remainder of this chapter the question, "Is the card a letter (number)?" is referred to as the "dominant question."

Subjects. The Ss were 144 students enrolled in the introductory psychology course at the University of Illinois.

Procedure. The nature of the task was explained to S who was then given two practice trials. After the second trial, S was asked to elaborate what he considered to be the best strategy. (See Appendix A for complete details).

Table 2

Design of Experiment I

8 letters	A	B	C	D	E	F	G	H	1	2	3	4	5	6	7	8
9 "	A	B	C	D	E	F	G	H	J	1	2	3	4	5	6	7
10 "	A	B	C	D	E	F	G	H	J	K	1	2	3	4	5	6
11 "	A	B	C	D	E	F	G	H	J	K	L	1	2	3	4	5
12 "	A	B	C	D	E	F	G	H	J	K	L	M	1	2	3	4
13 "	A	B	C	D	E	F	G	H	J	K	L	M	N	1	2	3
14 "	A	B	C	D	E	F	G	H	J	K	L	M	N	P	1	2
15 "	A	B	C	D	E	F	G	H	J	K	L	M	N	P	Q	1

Results

Two aspects of the data were analyzed. Examined first was the proportion of Ss who used the dominant question first. The χ^2 test for independent groups was used to test the hypothesis that the proportion was the same for the eight treatment groups. The value of χ^2 was significant ($\chi^2 = 57$, df = 7, $p < .001$) and the hypothesis was rejected. The proportions are shown in Table 8 of Appendix B and are plotted as a function of the number of letters in the source in Figure 2. Also shown is the average information or reduction in uncertainty (U) obtained with the dominant question. Both the proportion and the average information can vary between zero and one and are plotted on the same scale.

Examined second was the average value of $M(\sigma)$. The sample means and

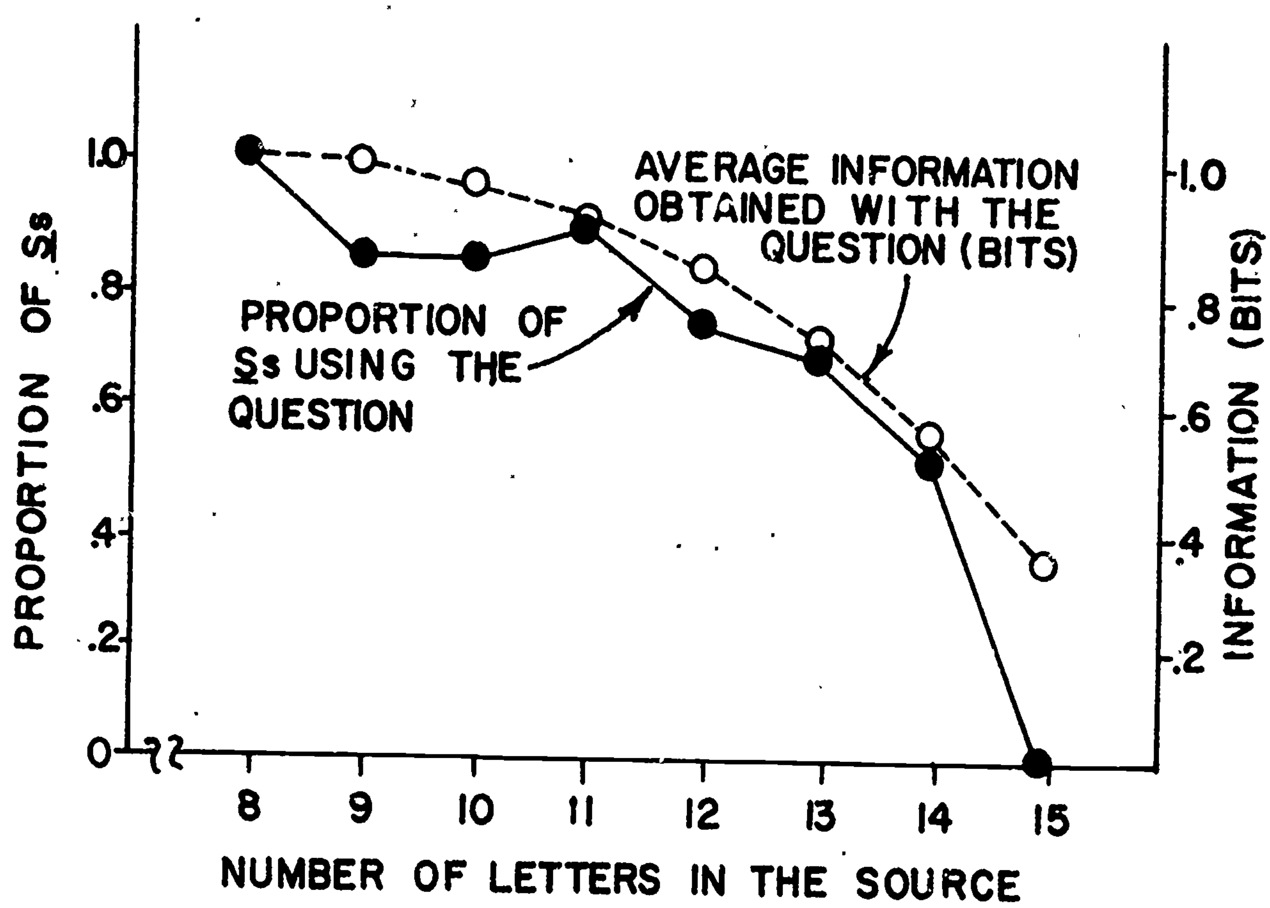

Figure 2. The proportion of Ss using the dominant question as a function of the number of letters in the source (U = 4, n = 16).

standard deviations are shown in Table 8 of Appendix B. Frequency plots of the data indicated that the population distributions were non-normal and were not homogeneous in terms of variance. Therefore, the non-parametric Kruskal-Wallis analysis of variance (Siegal, 1956) was used to test the hypothesis that the treatment populations were identical. The value of H was significant ($H = 36.1$, $df = 7$, $p < .001$) and the hypothesis was rejected. The average value of $M(\sigma)$ is plotted as a function of the number of letters in the source in Figure 3. Also shown as a reference are the values of $M(\sigma)$ that would have been obtained if S had divided the cards according to the letter-number distinction on the first question and had used optimal questions after that. This is the lowest value of $M(\sigma)$ which could be obtained when the dominant question was used first.

Discussion

The results indicate that Ss used the dominant question as long as it was not extremely inefficient (Figure 2). When the question became highly inefficient it was dropped in favor of other questions. It is interesting to note that when the question resulted in the maximum amount of information (one bit when there were eight letters), all Ss used it. When the question resulted in the minimum amount of information (about .34 bits when there were 15 letters), no Ss used it. At the extremes, then, Ss were in complete agreement regarding the use or non-use of the question.

The rate of change in the proportion of Ss using the question was small for relatively low inefficiencies (from eight to 11 letters) but increased as the inefficiency became larger. A comparison of the two curves in Figure 2 indicates that the proportion of Ss using the question was roughly proportional to the average information obtained with it.

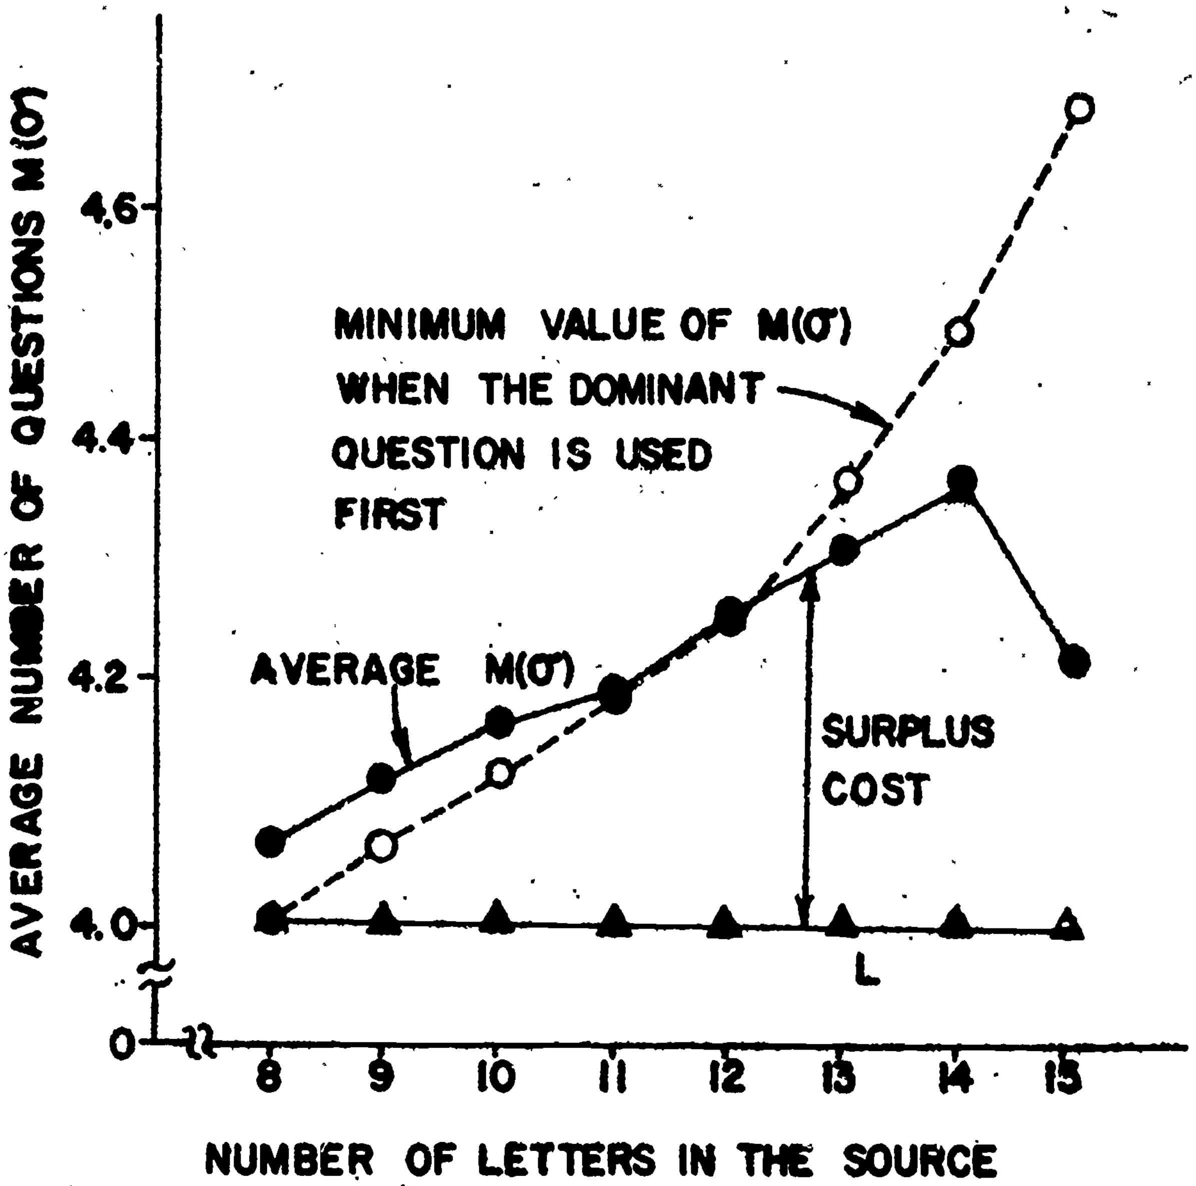

Figure 3. The average value of M(σ) as a function of the number of letters in the source.

This finding provides evidence that Ss were sensitive to the average amount of information which was obtained with the question. As long as the amount was not too low for the dominant question, it was likely to be used. However, as the average amount of information decreased, it was likely to be dropped.

The results shown in Figure 3 indicate that when Ss abandoned the dominant question, it was generally replaced with more efficient questions. If this were not the case, the average value of $M(\sigma)$ would have steadily increased as a function of the number of letters in the source. Instead, there actually was a large decrease in average $M(\sigma)$ when the number of letters changed from 14 to 15. Also, the average value of $M(\sigma)$ would have been above the value given by the reference line (the value of $M(\sigma)$ that would be obtained if S asked the dominant question first and was optimally efficient after that). The fact that the curve for the average value of $M(\sigma)$ dropped below the reference line as the number of letters increased is an indication that Ss were using more efficient questions in place of the dominant question.

The minimum average surplus cost occurred when there were eight letters in the source. In this case, Ss would have paid about six cents extra per trial if each question cost one dollar. Of the 20 Ss of this group, 13 employed a perfect strategy. The maximum surplus cost occurred when there were 14 letters. In this case, Ss would have paid about 38 cents extra per trial. When there were 15 letters in the source, the average surplus cost decreased to 22 cents per trial. If Ss had not dropped the dominant question in this case, the surplus cost would have been at least as high as 68 cents. In terms of cost then, Ss managed to keep their loss low relative to what it could have been.

Most Ss arranged the cards in two rows -- one of letters and one of numbers. Many of those who were optimally efficient arranged the cards into a 4 x 4 array and asked their questions in terms of rows and columns. A large number of Ss were aware that the dominant question was inefficient but used it anyway. Some insight into the reason for this can be obtained from a comment made by an S after the experiment: "I looked at it as being two problems. The first was to find the correct card when it was a letter and the second was to find the correct card when it was a number. I used my first question to determine which problem I was dealing with and then split the rest of the cards in two."

The results of this experiment show that Ss preferred a question which was based on a dominant characteristic of the stimuli as long as it was not highly inefficient. Their tendency to use such a question decreased as the average amount of information obtained with it decreased.

EXPERIMENT II: THE USE OF STRATEGIES AS A FUNCTION OF SOURCE CHARACTERISTICS

A source can be described in terms of two parameters--the uncertainty (U) and the number of categories or cards (n). Two experiments were conducted in which strategies were observed as a function of one of these parameters while the other was held constant. The quantity n/U is a measure of the skewness of the source distribution. When the distribution is uniform or least skewed, this quantity takes on relatively low values. As the distribution becomes more skewed, this quantity becomes larger. For example, when U equals two, the minimum value of n is four and occurs when the distribution is uniform. Higher values of n for the case of U equal to two imply greater deviation from a uniform distribution or greater skewness (e.g., see Table 3 for U=2, n=8).

Table 3

A Source with U=2, n=8

Card	A	B	C	D	E	F	G	H
Probability	1/2	1/4	1/8	1/16	1/64	1/64	1/64	1/64

Therefore, with U held constant, the source distribution becomes more skewed as n is increased, and with n held constant, the source distribution becomes more skewed as U is decreased.

It was expected that efficiency would decrease as the skewness of source distribution increased. The reason for this is that Ss tend to ask questions which divide the cards into groups of approximately equal number. If the source distribution is almost uniform, this type of question is relatively efficient. However, if the source distribution is highly skewed, this type of question is relatively inefficient.

Method

Procedure. In these experiments each card was labeled with a letter and the probability that it would be correct on any given trial. The letter appeared in the upper left-hand corner of the card and the probability was written in the center. During preliminary studies it was observed that some Ss did not make use of the source probability weights in developing their strategy. Instead the fractions indicating the probabilities were used only as a means of identifying the cards. For example, in the case of the experimental condition of Table 3, the question, "Does the correct card have 1/64 written on it?" is extremely inefficient and its use is an indication that S might be using the probabilities as labels.

Table 4

Control Condition for the Case of Table 3

Card	A	B	C	D	E	F	G	H
Label	A	B	C	D	E_1	E_2	E_3	E_4

Since they were not responding to the probabilities as relative frequencies, it was desired to eliminate these Ss from the examination of efficiency. In order to do this, Ss were first observed in a control condition in which each card was equally likely to be correct and in which the pattern of labels was identical to the pattern of probabilities in the experimental condition. For example, the control condition for the case of Table 3 is shown in Table 4. Cards with identical probabilities in the experimental condition became cards with identical letters but different subscripts in the control condition.

After elaborating a strategy for the control condition, S was observed in the experimental condition where the probabilities were written on the cards. If S asked the identical questions in terms of the pattern of labels for the control condition and in terms of the pattern of probabilities for the experimental condition, it was taken as evidence that the fractions were being used as labels only. If S responded differently to the two situations, it was taken as evidence that the fractions were being used to some extent as relative frequencies.

The complete procedure is given in Appendix A.

Table 5

The Conditions for Experiment II-A

n	Experimental									Control								
	A	B	C	D	E	F	G	H	I	A	B	C	D	E	F	G	H	I
6	1/2	1/4	1/16	1/16	1/16	1/16				A	B	C_1	C_2	C_3	C_4			
7	1/2	1/4	1/8	1/32	1/32	1/32	1/32			A	B	C	D_1	D_2	D_3	D_4		
8	1/2	1/4	1/8	1/16	1/64	1/64	1/64	1/64		A	B	C	D	E_1	E_2	E_3	E_4	
9	1/2	1/4	1/8	1/16	1/32	1/128	1/128	1/128	1/128	A	B	C	D	E	F_1	F_2	F_3	F_4

Subjects. Seven groups of 20 Ss were used. Each group consisted of ten males and ten females. The 140 Ss were taken from the groups of Experiment I.

Use of Strategies as a Function of n, with U held Constant--Experiment II-A

The four conditions of this experiment are shown in Table 5. For each of these conditions, U equals two bits, but n varies from 6 to 9.

Results

Two variables were examined as a function of n: (1) the proportion of Ss who used the fractions as labels; (2) the average value of M(O) for Ss who used the fractions as relative frequencies.

The hypothesis tested first was that the proportion of Ss who used the fractions as labels was the same for the four treatment populations. The expected frequencies for the four treatment groups under the null hypothesis were less than four. Therefore, in order to increase the expected frequencies the four groups were combined into two groups. The first group consisted of the group with n equal to six and the group with n equal to seven. The second group consisted of the group with n equal to eight and the group with n equal to nine. The value of χ^2 was significant ($\chi^2 = 5$, df = 1, $p < .05$) and the hypothesis was rejected. The proportion of Ss who used the fractions as labels is given in Table 9 of Appendix B.

The hypothesis tested second was that the population distributions of $M(\sigma)$ were identical for the different values of n. The Kruskal-Wallis analysis of variance was used for this purpose. The value of H was significant (H = 8.5, df = 3, $p < .05$) and the hypothesis was rejected. In Figure 4, the average value of $M(\sigma)$ is plotted as a function of n. The average values and standard deviations are given in Table 9 of Appendix B.

Discussion

It is clear that Ss approach to a situation depended in part on the number of cards in the source. As n increased, Ss were less likely to use identical strategies for the control and experimental conditions. (Appendix B, Table 9). This means that Ss were more likely to use the fractions as relative frequencies when n was large and the distribution was more skewed. Figure 4 shows that when Ss did make use of their knowledge of the probabilities, they were more efficient at lower values

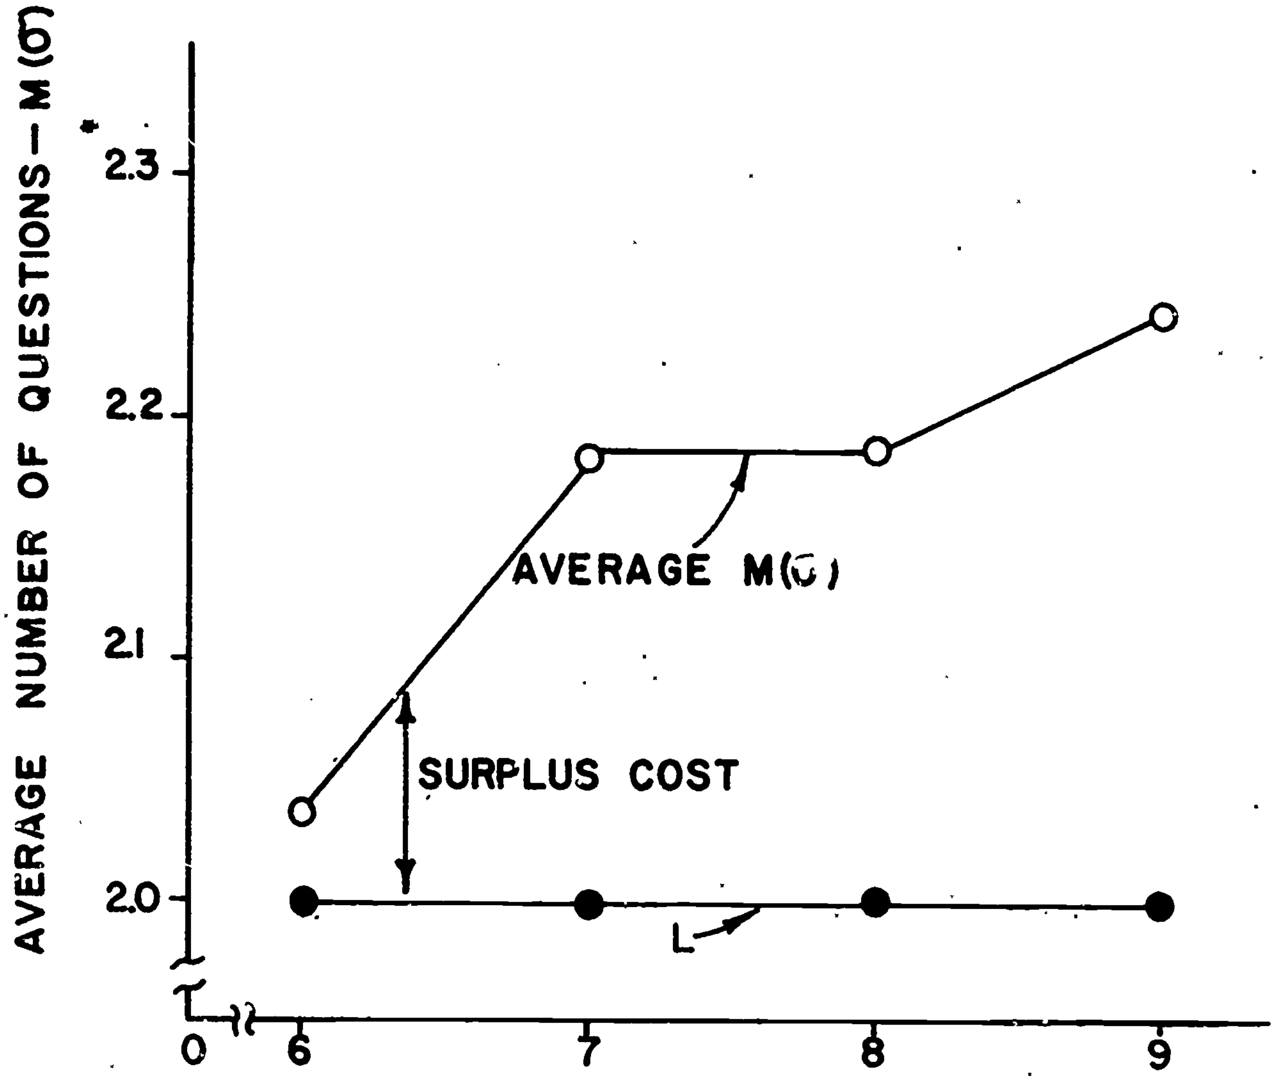

Figure 4. Average number of questions as a function of the number of cards in the source (U = 2).

of n. The surplus cost of Ss when n equalled six was less than .04 questions per trial. However, when n equalled nine it was over .24 questions per trial -- more than six times greater.

The reason for the decrease in efficiency as n increased is that Ss resisted dividing the cards into highly unbalanced groups in terms of number of cards. For n equal to six, the most efficient first question ("Is the card A?") divided the set into a group of one and a group of five cards. For n equal to nine the most efficient first question remained the same, but the cards were divided into a group of one and a group of eight -- a greater unbalance in terms of number of cards. In this situation Ss were more likely to ask questions such as, "Is it A or B?" and, "Is it A, B, or C?" These questions were less efficient but did not divide the set of cards into highly unbalanced groups.

The tendency of Ss to divide the cards into groups which are balanced in terms of n seems to be a result of a concern with the risk of using a large number of questions. Under any circumstance, a strategy which divides the cards into groups of equal n has a lower maximum possible number of questions than a strategy which does not. Therefore, if S were concerned with keeping the maximum possible number of questions low he should tend to ask questions which divide the set of cards into groups of equal size. Evidence that Ss were concerned with minimizing the maximum possible number of questions was also obtained in Experiment III and is discussed more fully there.

Use of Strategies as a Function of U, with n Held Constant--Experiment II-B

The four conditions of this experiment are shown in Table 6. For each of these conditions n equals eight, but U takes on the values 2, 2.25, 2.5, and 2.75. The condition in which U equalled two and n equalled eight was used in both Experiment II-A and Experiment II-B. Therefore, only one group of 20 Ss was observed in this condition and only seven groups were needed for both experiments.

Results

The analyses were identical to those of Experiment II-A except that the variables were examined as a function of U rather than n. The hypothesis tested first was that the proportion of Ss who used the probabilities as labels was the same for the four treatment populations. The value of χ^2 was significant ($\chi^2 = 8.125$, df = 3, $p < .05$) and the hypothesis was rejected. The proportion of Ss who used the probabilities as labels is shown for the different values of U in Table 10 of Appendix B.

Table 6

The Conditions for Experiment II-B

U	Experimental								Control							
	A	B	C	D	E	F	G	H	A	B	C	D	E	F	G	H
2	1/2	1/4	1/8	1/16	1/64	1/64	1/64	1/64	A	B	C	D	E_1	E_2	E_3	E_4
2.25	1/2	1/8	1/8	1/8	1/32	1/32	1/32	1/32	A	B_1	B_2	B_3	C_1	C_2	C_3	C_4
2.50	1/4	1/4	1/4	1/8	1/32	1/32	1/32	1/32	A_1	A_2	A_3	B	C_1	C_2	C_3	C_4
2.75	1/4	1/4	1/8	1/8	1/16	1/16	1/16	1/16	A_1	A_2	B_1	B_2	C_1	C_2	C_3	C_4

The hypothesis tested second was that the population distributions of $[M(\sigma) - L]$ were identical for the different values of U. The Kruskal-Wallis analysis of variance was used for this purpose. The value of H was significant ($H = 18$, df = 3, $p < .001$) and the hypothesis was rejected. The sample means and standard deviations of $M(\sigma)$ are shown in Table 10 of

Appendix B and the mean is plotted as a function of U in Figure 5.

Discussion

The results indicate that Ss approach to a situation depended in part on U. As U increased the Ss were generally more likely to use the fractions as relative frequencies (See Table 10, Appendix B). However, the proportion was not a monotonic function of U as was the case with n. Figure 5 shows that efficiency was greatest when U was greatest. In this case the average surplus cost was about .03 questions per trial. However, there was very little change in surplus cost when U varied from 2 to 2.5 bits -- the average surplus cost varied from .19 to .23 questions per trial.

It was expected that efficiency would increase as U increased and the skewness decreased. Although the results are in this direction, they are not as clear as was the case when n varied. A possible reason for this is that the fractions were used as both relative frequencies and labels. For example, when U equalled 2.5 there were three cards with a probability of 1/4. In order to ask an optimal question which divided the cards into groups of equal probability, this set of three cards had to be broken. However, 18 out of 20 Ss asked questions which grouped them together. The Ss did not break up the set because, as one said, "the cards seemed to go together." This is a further example of the tendency to ask questions which reflected the dominant characteristics of the stimuli. When U equalled 2.25 there was a similar set of three cards and Ss tended to group them together in their questions (14 out of 20). However, when U equalled 2 or 2.75, there were no groups of three identical probabilities and the pattern of probabilities was quite different. Therefore, the effect of U on efficiency was somewhat

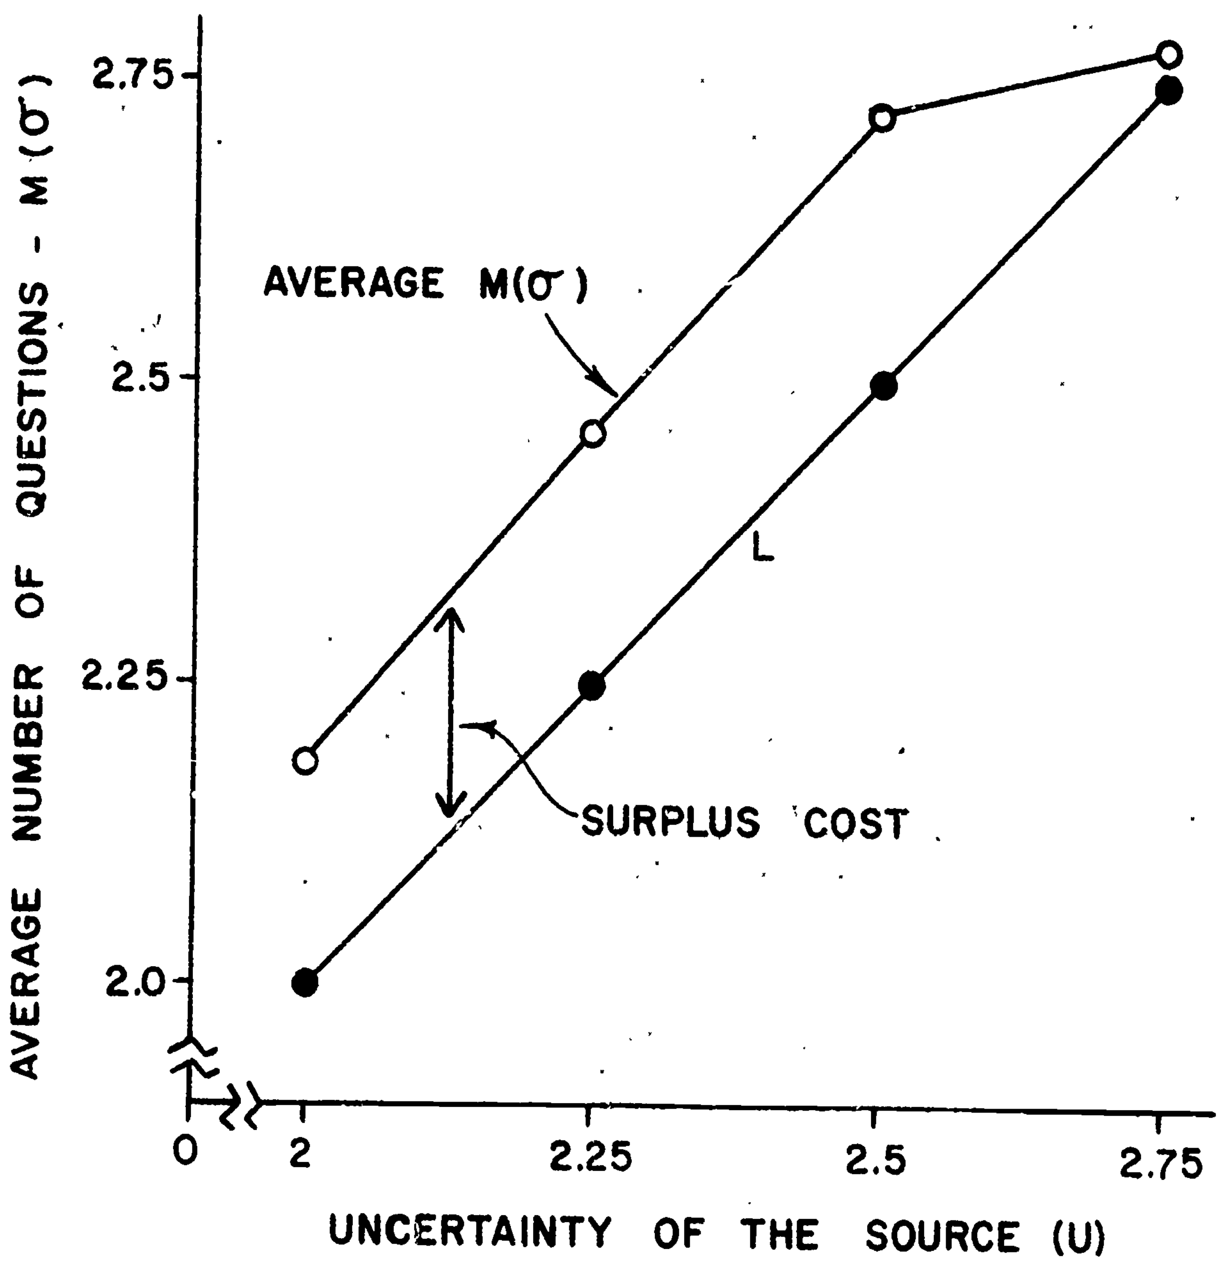

Figure 5. Average number of questions as a function of the source uncertainty ($n = 8$).

obscurred by the fact that the pattern of fractions varied and Ss questions were to some extent a reflection of the pattern.

When n varied, however, the pattern of fractions did not vary. Each situation had a series of decreasing probabilities and four identical probabilities at the low end. Therefore, changes in S's efficiency in this case could be attributed to the effect of n alone. The overall conclusion for experiments II-A and II-B is that Ss strategies were influenced by the skewness of the source distribution. As the skewness increased they were more likely to weight the fractions as relative frequencies as opposed to using them as labels only. Also, as the skewness increased those who did use the fractions as relative frequencies became less efficient. The decrease in efficiency is attributed to a concern with minimizing the maximum possible number of questions.

EXPERIMENT III: CHANGE OF STRATEGY AS A FUNCTION OF EXPERIENCE

Dale (1959) found evidence for improvement of strategies with experience, but his results were not conclusive, and Detambel and Stolurow (1957) observed that their Ss improved very little with experience. However, in those experiments, the Ss were not given feedback concerning how well they were doing. The purpose of this experiment was to examine changes of strategy as a function of experience when Ss were given feedback concerning their efficiency. It was expected that Ss who initially used optimal or near-optimal strategies would not be likely to change. On the other hand, Ss whose initial strategies were not optimal would be likely to change and to increase their efficiency.

Method

The 140 Ss from Experiment II were given ten trials with the source shown in Table 7. They were told that the best they could do in the long

Table 7

A Source with U=2, n=10

Card	A	B	C	D	E	F	G	H	I	J
Probability	1/2	1/4	1/8	1/16	1/32	1/64	1/256	1/256	1/256	1/256

run was to average two questions per trial. After each trial they were told the number of questions they used. Prior to and after the trials each S was asked to elaborate his best strategy (See Appendix A for details).

Results

The values of M(O) were calculated for the pre-experience strategies and for the post-experience strategies. The Ss were divided into two groups on the basis of the pre-experience M(O). The first group consisted of 43 Ss whose strategy was optimal or near-optimal ($2 \leqslant M(\sigma) \leqslant 2.005$). The second group consisted of the 97 Ss whose strategies were not optimal ($M(\sigma) > 2.005$). Separate analyses were performed on the groups.

In the optimal group, 9 out of 43 Ss changed strategy after experience (about 21%). For the non-optimal group, 81 out of 97 Ss changed strategy (about 83%). It is clear that Ss in the non-optimal group were likely to change strategy while those in the optimal group were not. Also, of the 81 Ss who changed strategy in the non-optimal group, 65, or 81%, decreased their value of $M(\sigma)$.

For the Ss in the non-optimal group $\Delta M(\sigma) = M_2(\sigma) - M_1(\sigma)$ was determined where $M_2(\sigma)$ was calculated from the post experience strategy and $M_1(\sigma)$ was calculated from the pre-experience strategy. The hypothesis that $\Delta M(\sigma) = o$ s tested against the alternative that $\Delta M(\sigma) \neq o$. The t test was used for this purpose. The value of t was significant ($t = -5.5$, $df = 96$, $p < .001$) and the hypothesis was rejected. The sample mean of $\Delta M(\sigma)$ was -.20 and the sample S.D. was .355. The average value of $M_1(\sigma)$ was 2.59 and the average value of $M_2(\sigma)$ was 2.39.

Discussion

It is clear that experience led to the development of efficient strategies. If Ss had to pay a dollar for each question, the experience would have resulted in a savings of about twenty cents per trial for those Ss who were initially inefficient. Before the experience they would have averaged about $2.59 per trial as opposed to $2.39 per trial after experience. Also, the fact that Ss who were initially efficient tended to remain so is an indication that efficient strategies persisted in the face of experience.

During the trials many Ss tried strategies which were actually more efficient than the one they finally chose. Also, the post-experience strategy for 25 Ss was less efficient than the pre-experience strategy. This apparent tendency toward inefficiency is a result of efforts to minimize the risk of using a large number of questions. The most efficient strategy for the case of Table 7 is one which goes from A to F, one card at a time, and then splits cards G, H, I, J into two groups of two cards. This strategy requires two questions on the average to determine the correct card. However, a very small percent of the time it will require as many as eight questions (when the correct card is G, H, I or J). Many Ss who were using a relatively efficient strategy became disconcerted when a low probability card was correct. On the following trials there was likely to be a change to a strategy which was somewhat less efficient but which would require at most five or six questions. For example, the strategy which groups cards two at a time starting with A and B requires an average of 2.3 questions. In this case, however, the maximum number of questions is only five.

Most Ss settled on a fixed strategy before the ten trials were completed. This is taken as evidence that even with extensive experience (say 100 trials) Ss would not become much more efficient. It is expected that the surplus cost would remain close to .39.

The conclusion is that Ss were able to improve efficiency as a result of experience. The improvement was limited due to a concern on the part of Ss with the maximum possible loss. In the attempt to keep this low there was a sacrifice of efficiency.

EXPERIMENT IV: TRANSFER OF LEARNING ACROSS SITUATIONS

When S develops a strategy for a particular situation, an interesting question is whether this reflects an approach which is specific to this situation or whether this reflects an approach that generalizes to different situations. Also of interest is whether the degree of generalization is sensitive to differences between the original situation and the new situation. The purpose of this experiment was to examine transfer of learning across situations.

If S is asked to elaborate his best strategy before and after experience in a particular situation, then $\Delta M(\sigma)$ is a reflection of a change in approach to the situation. When $\Delta M(\sigma)$ is negative, it indicates that S has become more efficient and is less concerned with the pattern of labels or with the risk of taking a great number of questions. When $\Delta M(\sigma)$ is positive, it indicates that S has become more concerned with minimizing the maximum number of questions than with being efficient. When $\Delta M(\sigma)$ is zero it indicates that S has not changed approach as a result of the experience.

Suppose that before and after experience on one situation S is asked to elaborate his best strategy for a <u>different</u> situation. The value of $\Delta M(\sigma)$ for the different or <u>transfer situation</u> could then be compared with that of $\Delta M(\sigma)$ for the situation on which experience is given (the <u>learning situation</u>). If S changes approach to the learning situation ($\Delta M(\sigma) \neq 0$), then a change in approach to the transfer situation could be a result of either generalization or uncontrolled, extraneous factors (e.g., S could change strategy just for the sake of a change). If S does not change approach to the learning situation

($\Delta M(0) = 0$), then a change in approach to the transfer situation could be a result of extraneous factors only. Therefore, generalization of approach is demonstrated when Ss who change approach to the learning situation are more likely to change approach to the transfer situation than Ss who do not change approach to the learning situation.

Generalization from one learning situation to several transfer situations can be examined by comparing values of $\Delta M(0)$ for the transfer situations. If the likelihood that Ss change strategy is different for the different transfer situations, it can be concluded that generalization of approach changes over the situations. The parameters n and U can be used as measures of similarity of situations. For instance, if U is the same for two situations, then similarity decreases as the difference in n increases. Likewise, if n is the same for two situations, then similarity decreases as the difference in U increases. By using these relationships, generalization can be examined as a function of similarity between the learning and transfer situations.

Procedure

The 140 Ss from Experiments II and III were observed again in their respective situations of Experiment II. For the purposes of this experiment, then, Ss elaborated strategies for two situations both prior to and after experience on one of them. For all Ss the learning situation was that of Experiment III (Table 7). The transfer situations were the treatments of Experiment II (Tables 5 and 6).

Results and Discussion

Two values of $\Delta M(0)$ were calculated for each S: one for the learning situation and one for the transfer situation. The Ss were divided into two groups on the basis of $\Delta M(0)$ for the learning situation. The first group consisted of the 90 Ss for whom $\Delta M(\sigma)$ was non-zero. The second group consisted of the 50 Ss for whom $\Delta M(\sigma)$ was zero. The proportion of Ss who changed approach to the transfer situation was compared for the two groups. When $\Delta M(\sigma)$ was zero for the learning situation, 11 out of 50 Ss changed approach in their transfer situation (about 22%). When $\Delta M(\sigma)$ was not zero for the learning situation, 52 out of 90 Ss changed approach in their transfer situation (about 58%). The difference in these proportions was significant ($\chi^2 = 16.6$, df = 1, $p < .001$). Of the 52 Ss who changed approach in both the learning and transfer situations, 39 changed in the same direction in both (75%).

These results provide evidence that the development of a strategy for one situation reflected an approach which generalized to other situations. The Ss who changed approach to the learning situation were more likely than those who did not to change approach to the transfer situation. The fact that only 58% of those who changed in the learning situation also changed in the transfer situation is a result of initially optimal performance by many Ss in the transfer situation. Most of these Ss improved efficiency in the learning situation but did not change their optimal approach to the transfer situation.

The fact that 75% of the Ss changed in the same direction for both situations is evidence that transfer tended to be positive. However, it is difficult to explain the 25% who became more efficient in one situation and less efficient in the other. Whether this was a result of negative transfer or extraneous factors is a question for further research.

Transfer as a function of source characteristics. The Ss who changed approach to the learning situation ($\Delta M(\sigma) \neq 0$) were classified according to their particular transfer situation. The proportion of $\underline{S}$s who changed approach in the same direction on both the learning and transfer situation was examined as a function of one of the source characteristics with the other held constant. Figure 7 shows this proportion plotted as a function of n with U held constant. The proportions were significantly different ($\chi^2 = 9.85$, $df = 3$, $p < .02$). Figure 8 shows the proportion plotted as a function of U with n held constant. The proportions were significantly different ($\chi^2 = 11.97$, $df = 3$, $p < .01$). The summary data are shown in Table 11 of Appendix B.

The results provide evidence for a gradient of generalization as a function of similarity between the learning and transfer situations. For the learning situation, U equals two and n equals ten. The sources of Figure 6 all have U equal to two, but n varies from six to nine. Therefore, the situation with n equal to nine is most similar to the learning situation, while the situation with n equal to six is least similar. It is evident that generalization decreased as similarity in terms of n decreased. The sources of Figure 7 all have n equal to eight but U varies from 2 to 2.75. Therefore, the situation with U equal to two is most similar to the learning situation while the situation with U equal to 2.75 is least similar. It is evident that generalization decreased as similarity in terms of U decreased.

One caution is necessary in regarding the results as evidence for a gradient of generalization as a function of similarity. It could be that under any circumstance $\underline{S}$s are less likely to change strategy for situations with lower n or higher U. If this were the case, one would expect to obtain results which are similar to those shown in Figures 6 and 7. In order to eliminate this alternative explanation, it is necessary to use a learning situation for which n is lower than those of the transfer situations or for which U is greater than those of the transfer situations.

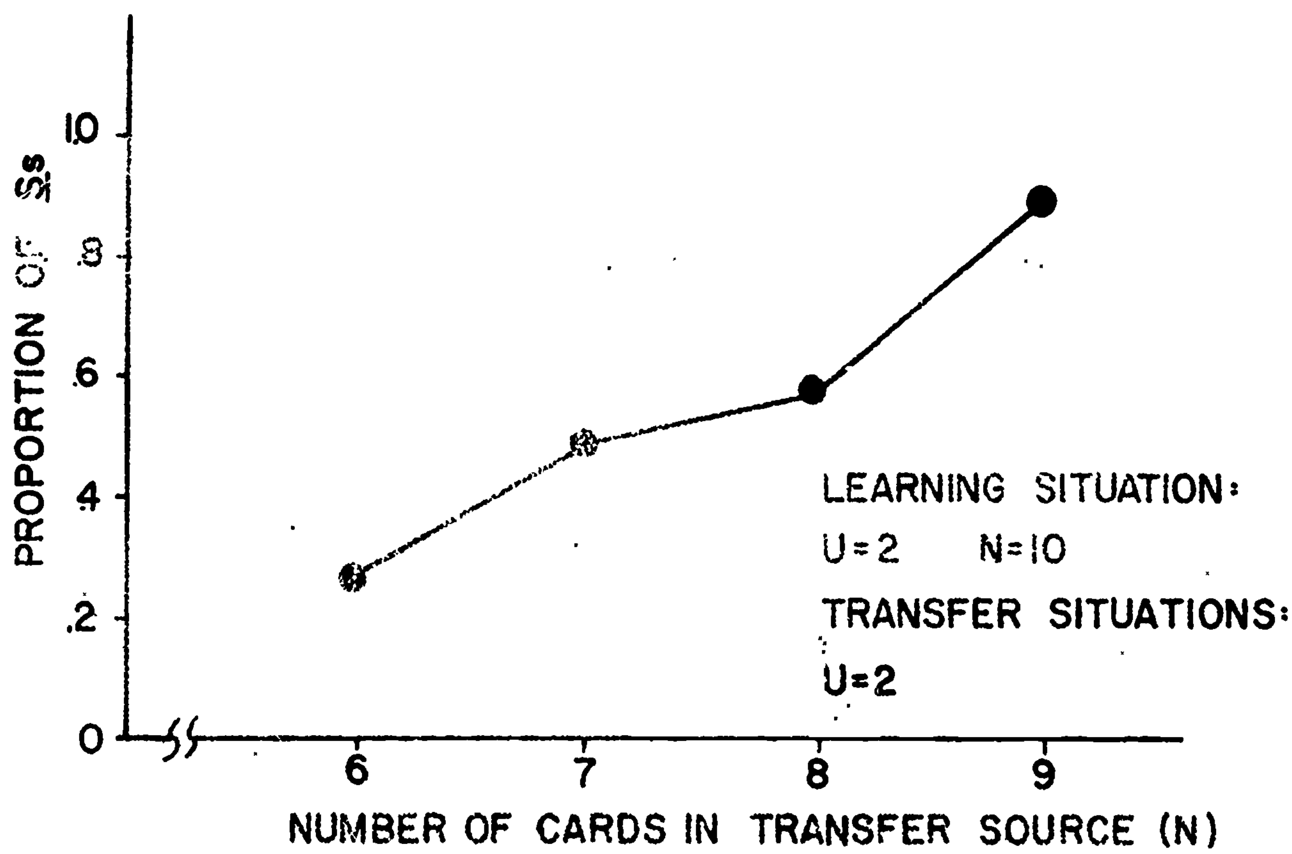

Figure 6. Proportion of Ss with $\Delta M(O) \neq 0$ in the learning situation who changed strategy for each of four transfer situations (n = 6, 7, 8 or 9).

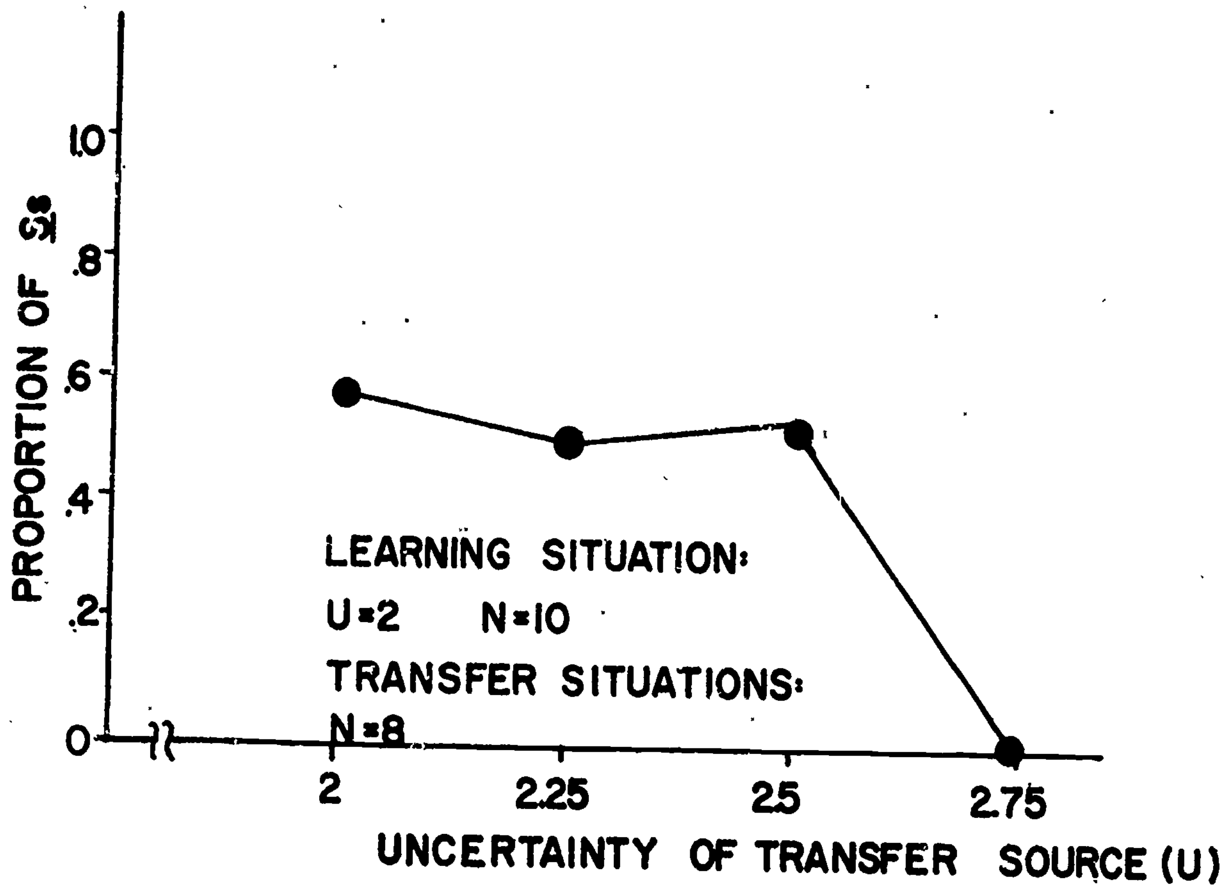

Figure 7. Proportion of Ss with $\Delta M(\sigma) \neq 0$ in the learning situation who changed strategy for each of four transfer situations (U = 2, 2.25, 2.5, or 2.75).

If there is a gradient of generalization in terms of similarity, the proportion of Ss who change would now decrease as n increases, and increase as U increases.

SUMMARY AND CONCLUSIONS

In general, the Ss were more efficient than was expected on the basis of previous results (Detambel and Stolurow, 1957; Dale, 1959). They were least efficient in the situation of Experiment III (U=2, n=10) prior to experience. In this case the average surplus cost was .59 questions per trial. Since L equalled two, the Ss used about 30% more questions than were absolutely necessary to determine the correct card. Approximately 1.3 questions were required on the average to obtain one bit of information when it was possible to average one question per bit.

In most cases, however, the inefficiency as calculated above was less than 10%, and dropped as low as 1%. For example, it was 1.1% when U equalled 2.75 bits in Experiment II-B. Here, approximately 1.01 questions were required on the average to obtain one bit of information. It is seen that the range of inefficiency was quite large -- 1% to 30%. This indicates that Ss were able to perform optimally, but also allowed themselves to become highly inefficient.

The overall results of the experiments indicate that three main factors influenced the use of strategies. They were: (1) the extent to which the questions reflected the dominant characteristics of the stimuli; (2) the average amount of information which was obtained with questions; (3) the risk of having to use a large number of questions. In Experiment I it was found that a question based on the dominant characteristics of the stimuli was preferred as long as it was not highly inefficient. When the question became highly inefficient, it was dropped in favor of less dominant but more efficient questions.

In Experiment II it was found that efficiency generally decreased as the skewness of the source distribution increased. This was a result of the tendency of Ss to ask questions which halved the cards in terms of number rather than probability. It was inferred from comments of Ss that this tendency resulted from a concern with the risk of asking a large number of questions.

In Experiment III, it was observed that experience generally led to the development of more efficient strategies. This indicates that the Ss were adaptive in the sense that they learned to pay less for each bit of information obtained from the environment. However, in some cases, efficiency was sacrificed to keep the risk of using a large number of questions low. In Experiment IV, a change in approach to one situation was reflected in a change to similar situations. This provides evidence that strategies involved the use of rules for applying the three factors which generalized across situations.

Various approaches to decision behavior concentrate on one of the three factors. The Expected Utility hypothesis (Edwards, 1961) views decision behavior in terms of maximizing efficiency or expected gain, while many applications of game theory are in terms of minimizing risk or maximum possible loss (Thrall, et. al., 1954). On the other hand, Simon (1956) proposes a model of choice behavior in which the primary factor is the structure of the environment.

In many simplified situations behavior is primarily a function of only one of the three factors. However, in the light of the results of this report and similar results obtained by Bruner, et. al. (1956), it is felt that complex decision behavior is determined on the basis of all three factors: efficiency, risk, and the structure of the environment.

A possible area for future work is the construction and evaluation of models describing strategy development in terms of weighting and combining the three factors. Such a model would take the form of a flow diagram in which Ss internal standards regarding the factors and the particular stimulus situation determine the outcome of a decision process - the outcome being Ss choice of strategy.

REFERENCES

1. Bendig, A. W. Twenty questions: An information analysis. Journal of experimental Psychology, 1953, 46, 345-348.

2. Bruner, J. S., Goodnow, J. J., & Austin, G. A. A Study of Thinking. New York: Wiley, 1956.

3. Dale, H. C. A. Strategies of searching in two simple systems. American Journal of Psychology, 1959, 72, 539-46.

4. Detambel, M. H. Probabilities of success and amounts of work in a multichoice situation. Journal of experimental Psychology, 1956, 51, 41-44.

5. Detambel, M. H., & Stolurow, L. M. Probability and work as determiners of multichoice behavior. Journal of experimental Psychology, 1957, 53, 73-81.

6. Edwards, W. Behavioral decision theory. In Annual Review of Psychology, 12, 473-98(Farnsworth, P. R., McNemar, O., McNemar, Q., Eds., Annual Reviews, Inc., Palo Alto, California, 1961.)

7. Garner, W. R. Uncertainty and Structure as Psychological Concepts. New York: Wiley, 1962.

8. Glaser, R., Damren, D., & Gardner, F. M. The tab item: A technique for the measurement of proficiency in diagnostic problem solving tasks. Educational and psychological Measurement, 1954, 14, 283-293.

9. Glaser, R., & Schwarz, P. A. Scoring problem-solving test items by measuring information. Educational and psychological Measurement, 1954, 14, 665-670.

10. Gluss, B. An optimum policy for detecting a fault in a complex system. Operations Research, 1959, 7, 468-77.

11. Goldbeck, R. A., Bernstein, B. B., Hillix, W. A., and Marx, M. H. Application of the half-split technique to a problem-solving task. Journal of experimental Psychology, 1957, 53, 330-38.

12. Posner, P. I. Information reduction in the analysis of sequential tasks. Psychological Review, 1964, 71, 491-504.

13. Rimoldi, H. J. A. Rationale and application of the test of diagnostic skills. Chicago: Loyola Univer., Psychometric Lab., Pub. No. 29, 1963.

14. Shannon, C. E., and Weaver, W. The Mathematical Theory of Communication. Urbana: University of Illinois Press, 1949.

15. Siegel, S. Nonparametric Statistics for the Behavioral Sciences. New York: McGraw Hill, 1956.

16. Simon, H. A. Rational choice and the structure of the environment. Psychological Review, 1956, 63, 129-38.

17. Stolurow, L. M., Bergum, B., Hodgson, J., and Silva, J. The efficient course of action in "trouble shooting" as a joint function of probability and cost. Educational and psychological Measurement, 1955, 15, 462-77.

18. Thrall, R. M., Coombs, C. H., and Davis, R. L., Eds., Decision Processes. New York: Wiley, 1954.